AFRL-OSR-VA-TR-2014-0335

INFORMATION DYNAMICS AS FOUNDATION FOR NETWORK MANAGEMENT

Mung Chiang
TRUSTEES OF PRINCETON UNIVERSITY

12/04/2014
Final Report

Air Force Research Laboratory
AF Office Of Scientific Research (AFOSR)/ RTC
Arlington, Virginia 22203
Air Force Materiel Command

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. REPORT DATE *(DD-MM-YYYY)*
11/30/2014

2. REPORT TYPE
Final Technical Report

3. DATES COVERED *(From - To)*
09/01/09 - 08/31/14

4. TITLE AND SUBTITLE
(MURI 09) Information Dynamics as Foundation for Network Management

5a. CONTRACT NUMBER

5b. GRANT NUMBER
FA9550-09-1-0643

5c. PROGRAM ELEMENT NUMBER

5d. PROJECT NUMBER

5e. TASK NUMBER

5f. WORK UNIT NUMBER

6. AUTHOR(S)
Chiang, Mung, Poor, Harold A., Rexford, Jennifer, (Princeton University)
Zhang, Junshan (Arizona State), Tropp, Joel (California Institute of Technology),
Candes, Emmanuel (Stanford University)
Carlson, Gunnar (Stanford University), Nowak, Robert (University of
Wisconsin-Madison), Markopoulou, Athina (University of California),

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
The Trustees of Princeton University
Office of Research and Project Administration
PO Box 36
Princeton, NJ 08544

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)
AFOSR/PKR - USAF, AFRL DUNS 143574726
AF Office of Scientific Research
875 North Randolph Street, RM 3112
Arlington, VA 22203
Alycia K. Armbruster - Alycia.Armbruster@us.af.mil

10. SPONSOR/MONITOR'S ACRONYM(S)
AFOSR/PKR

11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION / AVAILABILITY STATEMENT

Approved for public release; distribution is unlimited

13. SUPPLEMENTARY NOTES

14. ABSTRACT
An inter-disciplinary team (consisting of engineers, computer scientists, and mathematicians across foundation, theory and networked systems) made significant progress in understanding real-time information transmission over a variety of communication networks. Dozens of research breakthroughs have been made, ranging from geometry and topology to signal processing and communication protocol, from network optimization to systems implementation, and from security and privacy to radio technologies. The research findings have been widely disseminated through over 200 publications, including many best paper awards in top publication venues, and lead to multiple rounds of technology transfer.

15. SUBJECT TERMS
Communication, Networks, Wireless, Geometry, Topology, Optimization

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
					Mung Chiang
a. REPORT	**b. ABSTRACT**	**c. THIS PAGE**	U	9	**19b. TELEPHONE NUMBER** *(include area code)*
Unclassified	Unclassified	U			609-258-5071

AFOSR MURI 2009-2014
Information Dynamics as Foundation for Network Management

Final Report
November 2014

Summary:

The AF network is expected to support seamless transmissions of real-time information, including streaming video, voice and control messages, across disparate tactical data and IP networks, giving war fighters ubiquitous access to integrated shared information. Thanks to the support of this MURI project, a truly inter-disciplinary team (consisting of engineers, computer scientists, and mathematicians across foundation, theory and networked systems) made significant progress in understanding real-time information transmission. Multiple assistant professors were coached through projects funded by this MURI. Dozens of research breakthroughs have been made, ranging from geometry and topology to signal processing and communication protocol, from network optimization to systems implementation, and from security and privacy to radio technologies. The research findings have been widely disseminated through over 200 publications, including many best paper awards in top publication venues. Through close collaboration with AFOSR, AFRL, and multiple defense contractors, at least 5 cases of technology transfer have taken place successfully.

Publications:

Please see separate list.

Contribution Highlights:

Sheaf-theoretic Network Coding: We invented sheaf-theoretic tools for describing and classifying obstructions to linear network information coding. This gave a new, topological proof for weak duality in the fundamental theorem of network coding, and opened a new field of applications of sheaf theory to networks.

Sheaf-theoretic Max-Flow-Min-Cut Theorem: Motivated by the above breakthrough, We gave an entirely novel proof of the classical max-flow-min-cut theorem that relied on sheaf theory. This entailed the creation of a host of novel constructs: cellular sheaves taking values in semimodules, direct sheaf cohomology and cosheaf homology, and a version of Poincare duality adapted to sheaves of semimodules. These tools, though abstract, have led to a breathtakingly general version of flow-cut duality whose impact is poised to be very broad.

Sheaves and Sampling: Motivated by the above, we gave a novel interpretation of the classic Nyquist-Shannon sampling criterion for signals based on cellular sheaves. This opens new doors to obtaining sampling criteria without the use of harmonic analytic bounds.

Localization and Mapping via Signal Embeddings: Motivated by minimal-sensing scenarios with a general lack of localization (as exemplified in underground/indoor receivers, or adversarial/covert operations), one wants tools for extracting location or mapping an environment using passive opportunistic signals. The general framework for such a problem includes the complexities of multiple and perhaps anonymous transmitters, signal bounce, and signal modulation/drop-off. Ghrist and Robinson proved a very general signals embedding theorem, based on a modern update of the classic Whitney Embedding Theorem. The signals embedding theorem gives precise lower bounds on the "depth," or number of signals audible at each point in a domain, required to ensure unique channel response. This allows, e.g., the use of persistent homology in the signals space itself to reconstruct the topology of the domain from a collection of distributed non-collaborative receivers. These results have been verified experimentally using coarse acoustic data, ringtones on cell phones, and WLAN data from networks in building interiors.

Sparse modeling. We have proposed and analyzed new optimization formulations for recovering structured signals from an observation of their sum. We have established the first rigorous phase transition results for these models, which predict the exact situations in which the model is identifiable. These results depend on foundational research in conic geometry.

Random matrices. We have developed foundational results in random matrix theory and prepared expository material on these results. These techniques provide a simple and effective way to obtain information about non-classical random matrices that arise in applications. This work has already had a significant impact in the fields of algorithms, combinatorics, statistics, signal processing, scientific computing, and other areas.

Fast dimension reduction. We have developed an influential analysis of fast dimension reduction techniques that can be used to accelerate large-scale linear algebra. These methods can also be applied in compressed sensing systems.

Robust statistics. We have proposed and analyzed several new techniques for linear modeling in the presence of outliers. In contrast with classical approaches, these methods are based on convex optimization, so they admit tractable algorithms and performance guarantees.

Applied optimization. We have developed optimization methods for large-scale problems that arise in statistics and learning. In particular, we proposed and analyzed a novel technique for nonnegative matrix approximation by linear programming. We have also formulated algorithms for fitting low-rank matrix models using max-norm regularization. These projects have had a significant impact in mathematical signal processing and contingent research areas. The MURI funding has played a crucial role in catalyzing this work.

Deadline-aware Scheduling with Adaptive Network Coding for Real-time Traffic. We have studied deadline-aware scheduling with adaptive network coding (NC) for real-time traffic over the downlink of a wireless network. To meet hard deadlines of real-time traffic, the block size for NC is adapted based on the remaining time to the deadline so as to strike a balance between maximizing the throughput and minimizing the risk that the entire block of coded packets may not be decodable by the deadline. This sequential block size adaptation problem is cast as a finite horizon Markov decision process. Since the optimal solution to such deadline scheduling problems requires to explicitly take into account the coupling between the deadlines and the stochastic characteristics of wireless systems, we focus on developing low-complexity MDP-based scheduling algorithms. One interesting finding is that the optimal block size and its corresponding action space monotonically decrease as the deadline approaches, and that the optimal block size is bounded by the “greedy” block size. These unique structures make it possible to significantly narrow down the search space of dynamic programming, building on which we develop a monotonicity-based backward induction algorithm (MBIA) that can find the *optimal* block size in *polynomial time*. Further, a joint real-time scheduling and channel learning scheme with adaptive NC is developed to adapt to channel dynamics in a mobile network environment. We devise a low-complexity online scheduling algorithm integrated with the MBIA, and then establish its asymptotical utility optimality. The analysis and simulation results are corroborated by high fidelity wireless emulation tests, where actual radio transmissions over emulated channels are performed to demonstrate the feasibility of the MBIA in finding the optimal block size in real time.

Tech transfer. Our collaboration with Intelligency Automation Inc. has led to DoD technology transfer, and was recently highlighted by AFOSR upper administration. Specifically, the adaptive NC schemes are implemented and tested in a realistic wireless emulation environment with real radio transmissions. Our testbed platform consists of four main components: radio frequency network emulator simulator tool, RFnestTM, software simulator running higher-layer protocols on a PC host, configurable RF front-ends (RouterStation Pro from Ubiquiti), and digital switch. We removed the radio antennas and connected the radios with RF cables over an attenuator box. Then, real signals are sent over emulated channels, where actual physical-layer interactions occur between radios, and in the meantime the physical channel attenuation is digitally controlled according to the simulation model or recorded field test scenarios can be

replayed accordingly.

Diffusion of Real-Time Information in overlaying Social-Physical Networks. We have studied the diffusion behaviors of real-time information flow in an overlaying social-physical network. Typically, real-time information is valuable only for a limited time duration and hence needs to be delivered before its “deadline”. Capitalizing on the theory of inhomogeneous random graph, we analytically quantify the size of real-time information epidemic in a coupled social-physical network. One interesting finding is that a larger-size online social network may not necessarily yield a larger size of information epidemic. In fact, under certain conditions, the size of information epidemic could even decrease with the growing size of the online social network, and this is in stark contrast to the phenomena in a single network.

Throughput maximization via PHY-aware distributed scheduling: In tackling the challenges in network dynamics, probing/sensing is a key mechanism to estimate and track network/channel states for state-aware adaptive resource allocation. However, incorporation of stochastic dynamics into network optimization is often challenging, because of the intrinsic tradeoffs between probing cost and scheduling. In this project, we have investigated the fundamental tradeoffs between probing and scheduling for throughput maximization, in the context of a contention-based ad-hoc network. In such a network, channel-aware distributed scheduling involves a sequence of contention/probing and packet transmissions, and the desired tradeoff boils down to judiciously choosing the condition for stopping probing. Capitalizing on optimal stopping theory, we show that the optimal scheduling policy turns out to be a pure threshold policy. We generalize the study to cases with noisy probing, and then investigate distributed scheduling with two-level probing, aiming to quantify whether it is worthwhile for the link with successful contention to perform further probing to mitigate estimation errors, at the cost of additional probing. We have also made good progress in devising distributed scheduling under delay constraints.

Delay Analysis of wireless scheduling: Delay is another equally important performance metric, but remains under-explored due to technical difficulty. In this project, we first leveraged heavy traffic analysis for studying session-level dynamics. In particular, we carried out diffusion approximation for delay performance of wireless scheduling, where signaling overhead is accounted for by vacation models, and studied stability-delay-complexity tradeoff. Roughly speaking, heavy traffic approximation can be viewed as a Functional Central Limit Theorem, and is applicable to dynamic networks with

bottleneck links. Given that delay performance has particularly important implications in cognitive radio networks, we developed a stochastic fluid queue framework for analyzing the delay of secondary users in cognitive radio networks with random access.

Cooperative wireless networks: In earlier work we developed a new cooperative relaying strategy, that was shown theoretically shown to be approximately optimal for arbitrary wireless networks. In this project we built on this to demonstrate its instantiation through lattices as well as through iterative codes. This was also implemented on a wireless testbed. We had several other results on wireless networks, including results on interference management with bursty sources and resources, full duplex networks etc. These lines of work resulted in 8 publications in top-tier conferences and journals.

Wireless network secrecy: We have developed a new way of creating wireless network secrecy through intentional noise-insertion to create the right environment for interactive secrecy. This enabled us to not only prove new conclusive information-theoretic security results, but it was also demonstrated through an implementation on a wireless testbed. This line of work resulted in 8 publications in top-tier venues.

Source-channel computation and communication: One of the central results in information theory is the architectural result of source-channel separation by Shannon. It is known that the separation is not true in networks. However, we demonstrated that for many situations separation of communication and data compression/computation is either optimal or approximately optimal. This demonstrates that the gain of complicated joint techniques could be limited in many important situations. We demonstrated an architectural result on general networks without solving the component open questions, a technique that might have implications on other problems. This line of work resulted in two journal publications in top-tier venues.

Wireless coded caching and network coding: We have initiated a new study of jointly using broadcasting along with caching for content delivery. In recent work, we have demonstrated the approximate optimality of network coding techniques for this problem. We have also studied subspace properties of (non-coherent) network coding and its applications. These lines of work has resulted in 5 publications in top-tier venues.

Diagnosing Network Performance Problems for Data Center Applications

Network performance problems are notoriously tricky to diagnose, and this is magnified when (i) applications are often split into multiple tiers of application components spread across thousands of servers in a data center and (ii) the end-to-end traffic traverses multiple independently-administered networks. In this project, we have designed, prototyped, and deployed techniques for monitoring network performance on end-hosts and combining measurements from multiple locations to pinpoint the root cause of performance problems. We have developed monitoring techniques for both the

public cloud (where the cloud provider can instrument the server operating system) and the private cloud (where the cloud provider can monitor traffic in the hypervisor). Our results include:

- SNAP (scalable network application profiler), deployed in Microsoft's Bing service: Performance problems often arise in the communication between different parts of the distributed cloud application. Working with collaborators at Microsoft Bing, we designed SNAP, a scalable network-application profiler that guides developers in identifying and fixing performance problems. SNAP passively collects TCP statistics and socket-call logs with low computation and storage overhead, and correlates across shared resources (e.g., host, link, switch) and connections to pinpoint the location of the problem (e.g., send buffer mismanagement, TCP/application conflicts, application-generated microbursts, or network congestion). Our one-week deployment of SNAP in one of Microsoft's production data centers (with over 8,000 servers and over 700 application components) has already helped developers uncover 15 major performance problems in application software, the network stack on the server, and the underlying network. A paper on this work was published at the Networked Systems Design and Implementation (NSDI) Conference in 2011.

- Wide-area performance diagnosis, deployed in CoralCDN: We ported the SNAP tool to Linux, and used similar techniques to diagnose wide-area performance problems in content-distribution networks. By correlating performance across TCP connections destined to related groups of Web clients, our system can identify which network along the path is to blame for performance problems, or whether long round-trip times were to blame. We could also identify when limitations of the client host (e.g., limited buffer space) were responsible for performance problems. We deployed our system in the PlanetLab testbed and used it to diagnose performance problems on CoralCDN, an widely-used content distribution network. A paper on this work was published in the ACM SIGCOMM Workshop on Measurements Up the Stack in 2011.

- HONE programmable platform for private clouds: Building on our experience with SNAP, we designed a programmable platform for cloud providers to collect and analyze measurements of end hosts and network switches. Our HONE (HOst-NEtwork) system presents a uniform view of a diverse collection of measurement data, minimizes measurement overhead by performing lazy materialization of fine-grained statistics, and scales the analysis by processing data locally on the end hosts. HONE offers a simple and expressive programming framework for network and service administrators. We evaluated HONE by implementing several canonical traffic-management applications, measuring its efficiency with micro-benchmarks, and demonstrating its scalability with larger experiments on Amazon EC2. A paper on this work has been accepted for the Journal on Network and Systems Management in 2014.

- RINC programmable platform for Infrastructure-as-a-Service public clouds: SNAP and HONE require access to the network stack in the end-host operating system. While

appropriate in private clouds (e.g., Bing) and content-distribution networks (e.g., CoralCDN), this approach is not acceptable in public cloud settings where tenants run their own virtual machines (VMs). In our recent research on RINC, we collect TCP-level measurements in the hypervisor, rather than the network stack. While operating transparently to the tenants, RINC must now infer the TCP statistics from the observed packet stream. We design techniques that track the evolution of a TCP connection, including novel approaches for monitoring connections already in progress. RINC's ability to monitor in-progress connections enables us to perform light-weight monitoring of network traffic, and only collect detailed statistics on a subset of connections to (i) cover all applications, server machines, and network devices or (ii) reactively enable heavier-weight monitoring after detecting a performance problem. Experiments with our prototype system demonstrate that RINC enables IaaS providers to diagnoses network performance problems for their tenants at reasonable overhead. A paper on this work is in submission to the Symposium on Cloud Computing, 2014.

Flexible Network Policies Using Commodity Switch Hardware

Software-Defined Networking (SDN) enables a wide-range of innovative network services, such as flexible traffic monitoring, access control, server load balancing, network virtualization, and traffic engineering. These services rely on the ability to install a relatively large number of fine-grained packet-forwarding rules in the switches. Each rule matches on fields in the packet header, performs a limited set of actions (e.g., drop, forward), and counts the matching packets. However, commodity SDN switches support only a small number of rules (e.g., several thousand), because of the cost and power requirements of high-speed TCAM (Ternary Content Addressable Memory). In this project, we designed a collection of algorithmic techniques for optimizing the use of rule-table space, including:

- Incremental consistent updates: A consistent update installs a new packet-forwarding policy across the switches of an SDN in place of an old policy. While doing so, such an update guarantees that every packet entering the network either obeys the old policy or the new one, but not some combination of the two. While consistent updates ensure the network obeys invariants (e.g., no loops, no blackholes) during the transition, the technique can double the number of rules in the switches, in the worst case. In this work, we introduce new algorithms that trade the time required to perform a consistent update against the rule-space overhead required to implement it. We break an update in to k rounds that each transfer part of the traffic to the new configuration. The more rounds used, the slower the update, but the smaller the rule-space overhead. To ensure consistency, our algorithm analyzes the dependencies between rules in the old and new policies to determine which rules to add and remove on each round.
A paper on this work appeared that the HotSDN workshop in August 2013.

- Distributing rules over multiple switches: Many SDN applications, such as access

control and server load balancing, treat the network as one big virtual switch. The SDN controller takes responsibility for installing the resulting rules in the underlying physical switches. The simplest approach is to place the fine-grained rules that drop, modify, or forward traffic at the ingress switches. However, these switches can easily run out of rule-table space. In this work, we show how to distribute the rules over multiple switches along a path, while still obeying the "one big switch" policy as well as whatever routing the controller wants to enforce across the switching fabric. Our optimization technique works with arbitrary network topologies and routing policies, and switches with heterogeneous rule-table sizes. We also
designed incremental algorithms that minimize the updates to the switches when the policy changes.
A paper on this work appeared at the CoNEXT conference in December 2013.

- Incremental updates in a compositional hypervisor: SDN controller applications should be written in a modular fashion, with independently-written modules for (say) routing, monitoring, access control, and server load balancing. The controller can compose the results of these applications to generate a single set of rules to install in the underlying switches. However, each of these application modules change their policies over time in response to changes in network conditions. Recomputing and installing the resulting switch-level rules from scratch would consume substantial processing resources on the controller, and lead to disruptions in the underlying network. Instead, we have designed incremental algorithms for computing the changes in the rules, by recognizing that composition operators (like parallel and sequential composition) correspond to a convenient algebra on the rule priorities. We show that, through simple arithmetic on the rule priorities, the controller can easily determine which rules need to be added, changed, or removed, without altering any of the other rules. A paper on this work will appear at the HotSDN workshop in August 2014, and an extended version will be submitted in the early fall.

- Rule caching: Many rules in a network policy match a relatively small portion of the traffic. If we view the rule table in the switch as a cache, we can design algorithms that store the most popular rules in the cache and handle cache misses in software. However, we cannot simply apply conventional cache-replacement algorithms (like least recently used) because the rules have overlapping patterns. These overlapping patterns mean that removing one unpopular rule from the cache can cause data packets to incorrectly match a different rule, leading to incorrect behavior. We designs a compact data structure for maintaining the dependencies between rules, and design caching algorithms that respect these dependencies. One simple solution moves entire chains of dependent rules in and out of the cache, at the expense of wasting rule-table space on some unpopular rules. A second, more sophisticated algorithm "splices" these long dependency chains by generating new rules that "cover" a large number of unpopular rules, leaving more cache-table space for the popular rules. We have evaluated our solution on commercial SDN switches and are in discussions with Broadcom about applying these techniques in their switches. A paper

on this work will appear at the HotSDN workshop in August 2014, and an extended version is in submission to the CoNEXT Conference, 2014.

- Load balancing on commodity switches: Networks often need to divide traffic over multiple paths, middleboxes, or servers, based on a load-balancing policy. For example, a replicated Web site might divide one-third of its requests over each of three backend servers. Rather than implementing load balancing in dedicated appliances, commodity SDN switches can perform this function. We design algorithms that generate switch-level rules for dividing traffic over different components, given target load-balancing weights. When a network has many Web sites (each with multiple backend servers) or destinations (each with multiple network paths), the rule-table space quickly becomes a limitation. In this work, we design algorithms for (i) approximating the weights for each service, (ii) truncating each approximation to fit within a limited number of rules, (iii) packing the rules for different services based on their relative popularities, and (iv) clustering services with similar weights to share a common set of rules. Our algorithms enable commodity switches to implement load-balancing policies for large cloud providers. This is joint work with Google's networking team, and a paper is in submission to the CoNEXT conference, 2014.

CROSS LAYER OPTIMIZATION OF CODED WIRELESS NETWORKS

Wireless environments lend themselves to network coding, thanks to the broadcast nature of the wireless medium. We consider unicast flows over wireless networks with constructive inter-session coding schemes, including but not limited to COPE (one-hop opportunistic coding) and butterfly packing (two-hop coding). Such coding schemes are simple to implement and have been shown to significantly increase throughput. However, network coding-agnostic TCP and UDP flows cannot fully exploit the network coding opportunities. In this part of the project, we jointly optimized network coding, transport and application protocols.

NETWORK CODING AND INFERENCE

There is a close relation between topology and network coding: coding at intermediate nodes introduces topology-dependent correlation in the content of the packets, which can then be exploited to infer the topology. We developed a number of active probing techniques and specialized network coding schemes that allow inference of the topology based on end-to-end probes.

IMPLEMENTATION OF NETWORK CODING ON ANDROID PHONES

As part of this MURI, we started a thread on implementation of network coding on Android smartphones. This is a non-trivial task in general, and on mobile devices in particular, due to their limited resources. As part of this thread, we built several **efficient implementations of network coding libraries** on android, both in Java and one in C++. In addition, and perhaps even more importantly, we identified two application

scenarios and utilized these libraries as a building block to solving a bigger problem. More specifically, we designed two novel cooperative schemes: (1) Microcast - for video streaming and (2) Microplay - for local multiplayer games, both of which can benefit from the use of network coding in some of their components.

Microcast: Cooperative Video Streaming. Mobile video is one of the increasingly popular, as well as demanding, applications on smartphones. We consider a group of smartphone users, within proximity of each other, who are interested in streaming the same video from the Internet at the same time. The common practice today is that each user downloads the video independently using one connection (e.g., cellular or WiFi), which often leads to poor quality.

We designed, implemented, and evaluated a novel system, called MicroCast, which cooperatively uses the resources on all smartphones in the group so as to improve the streaming experience. Each phone uses simultaneously two network interfaces: one (cellular) to connect to the video server and another (WiFi) to connect to the rest of the group. Key ingredients of our design include the following. First, we propose a scheduling algorithm, MicroDownload, that decides what parts of the video each phone should download from the server, based on the phones' download rates and the congestion in the local network. Second, we propose a novel all-to-all local dissemination scheme, MicroNC-P2, for sharing content among group members, which outperforms state-of-the-art peer-to-peer schemes in our setting. MicroNC-P2 is explicitly designed to exploit WiFi overhearing and network coding, based on a local broadcast framework, MicroBroadcast, which we developed specifically for Android phones.

This work led to a spin-off from UC Irvine and EPFL, www.shoelacewireless.com, see Commercialization section later.

Microplay: Local Multiplayer Games: Smartphones are an ideal platform for local multiplayer games, thanks to their computational and networking capabilities as well as their popularity and portability. However, existing game engines do not exploit the locality of players to improve game latency. In this work, we designed MicroPlay, a complete networking framework for local multiplayer mobile games. To the best of our knowledge, this is the first framework that exploits local connections between smartphones, and in particular, the broadcast nature of the wireless medium, to provide smooth, accurate rendering of all players with two desired properties. First, it performs direct-input rendering (i.e., without any inter- or extrapolation of game state) for all players. Second, it provides very low game latency. We implement a MicroPlay prototype on Android phones, as well as an example multiplayer car racing game, called Racer, in order to demonstrate MicroPlay's capabilities. Our experiments show that cars can be rendered smoothly, without any prediction of state, and with only 20–30 ms game latency.

Online Social Networks Sampling: http://odysseas.calit2.uci.edu/osn

In this part of the project, we were interested in measuring large graphs, in particular in developing methods for sampling and estimation of various properties of interest. Our primary application domain and motivation was Online Social Networks (OSNs), and the methods and models have been optimized for this context accordingly.

Because sampling directly from the user population is often expensive (or impossible), link-trace (a.k.a. crawling) techniques are typically used. Early measurement studies, used graph exploration techniques, such as breadth-first-search (BFS), which do not necessarily provide a representative sample. We showed how to collect a probability sample of users by crawling the OSN in a principled way, rooted at the Monte-Carlo Markov-Chain literature. We presented a framework based on crawling of the social graph, and we obtained and analyzed the first uniform sample of Facebook users. We developed a multi-graph technique that allows to crawl multiple relations (e.g., friendship, co-membership in a group, participation in the same event) when the social graph is poorly connected or even disconnected. We demonstrated the efficiency of this approach in Last.FM, which is a fragmented network. We considered user attributes in addition to network structure and we showed how to perform stratified sampling via crawling, which is needed when not all nodes are equally important in the estimation objective. We designed a method for correcting the bias of BFS sampling. We showed that this method is exact in random graphs with a given degree distribution, and an efficient heuristic in real graphs.

We then switched our attention from sampling to generative models that can produce graphs that resemble the original ones in terms of various known properties (e.g. estimated through crawling). We developed and evaluated methods for graph estimation and construction that achieve different tradeoffs between complexity and accuracy. A core question is to identify the minimum set of features that need to be included in the model, so as to accurately capture characteristics of real graphs, while maintaining low complexity. We adopted the systematic framework of dk-series for characterizing the properties of a graph using a series of probability distributions specifying all degree correlations within *d*-sized subgraphs of a given graph G. Increasing values of *d* capture progressively more properties of G at the cost of higher complexity. Methods for deterministic graph construction and sampling from *dk*-series have been developed, albeit not always computationally efficient, up to *2K*, while only MCMC techniques are known for $d>2$. We developed and evaluated methods for 2K and 2.5K graph estimation and construction. By "2.5K", we define graphs that have a target 2K (i.e., joint degree distribution) and some notion of clustering (i.e., a given degree-dependent clustering coefficient or average clustering coefficient). This is important for social networks that exhibit high degree of clustering, thus cannot be modeled using *2K* alone. Our work is the first that can sample large online social networks and then construct synthetic graphs that resemble the original ones, in terms of the metrics defined above, on the order of a few hours. Our methodology can potentially be applied to network data beyond social networks as well. While the 2.5K construction algorithms were heuristic, in our recent work we designed a linear-time algorithm that can provably

construct graphs with a given joint degree distribution and provides more flexibility than prior methods.

All the datasets collected in the aforementioned studies have been made publicly available at http://odysseas.calit2.uci.edu/osn. They have been downloaded and used by more than 3000 researchers as of Aug. 2014. We have also made publicly available the software developed as part of this research.

Compressive sensing and signal processing.

Yuejie Chi (now an Assistant Professor of Electrical Engineering at Ohio State) received the 2013 IEEE Signal Processing Society Young Author Best Paper Award for

Y. Chi, L. Scharf, A. Pezeshki and R. Calderbank, Sensitivity to Basis Mismatch in Compressed Sensing, *IEEE Transactions on Signal Processing,* Vol. 59 (5), pp. 2182-2195, May 2011

We also considered the problem of reconstructing a data stream from a small subset of its entries, where the data stream is assumed to lie in a low-dimensional linear subspace, possibly corrupted by noise.

Y. Chi, Y.C. Eldar and R. Calderbank, PETRELS: Subspace estimation and tracking from partial observations, *Proceedings of the International Conference on Acoustics, Speech, and Signal Processing (ICASSP 2012)*. Kyoto, Japan, March 2012

was selected from more than 800 submissions to receive the *Google Student Paper Award* at ICASSP 2012.

Most recently she participated in the 2014 Summer Faculty Fellowship Program at AFRL, Rome NY.